

Text Annotation Graphs: Annotating Complex Natural Language Phenomena

Angus G. Forbes*, Kristine Lee[†], Gus Hahn-Powell[‡],
Marco A. Valenzuela-Escárcega[‡], Mihai Surdeanu[‡]

*University of California, Santa Cruz, angus@ucsc.edu

[†]University of Illinois at Chicago, khlee2@uic.edu

[‡]University of Arizona, {hahnpowell,marcov,msurdeanu}@email.arizona.edu

Abstract

This paper introduces a new web-based software tool for annotating text, *Text Annotation Graphs*, or *TAG*. It provides functionality for representing complex relationships between words and word phrases that are not available in other software tools, including the ability to define and visualize relationships between the relationships themselves (semantic hypergraphs). Additionally, we include a visualization mode in which annotation subgraphs, or semantic summaries, are used to show relationships outside of the sequential context of the text itself. These subgraphs can be used to quickly find similar structures within the current document or external annotated documents. *TAG* was initially developed to support information extraction tasks on a large database of biomedical articles. However, our software is flexible enough to support a wide range of annotation tasks for many domains. Examples are provided that showcase *TAG*'s capabilities on morphological parsing and event extraction tasks.

Keywords: annotation, event extraction, online NLP tools, text visualization




Figure 1: A screen capture of the annotation panel in the *TAG* software. Here a single sentence of biomedical text is shown, with both semantic (top) and syntactic (bottom) relationships toggled on.

1. Introduction

According to systems biologist Arthur Lander, the dominant icon of our current era is the “hairball” (Lander, 2010), a visual depiction of a network of such complexity that the edges overlap so densely as to obscure the meaning of the very relationships it is meant to explain. Part of the reason for this shift from a more straightforward symbol (such as the double helix, which, while still complex, indicates a more orderly, and potentially predictable machinery) to the hairball, Lander writes, is the rise of the computational tools that make it possible to generate vast amounts of “big data,” as well as the realization that thorny problems generally involve the intersections between data drawn from multiple, interacting systems.

Contemporary information extraction approaches seek to capture complex natural language phenomena, such as those found in biomedical literature (Björne et al., 2010; Valenzuela-Escárcega et al., 2015). These biomedical articles present new knowledge regarding biological processes which can be usefully described via a semantic graph that

represents the event structure of biochemical reactions. Moreover, these approaches can be used to detect predicates (triggers), relations, and events within individual articles in order to find connections to other corpora and to draw inferences about relationships between texts.

We introduce a novel web-based annotation software application, *Text Annotation Graphs*, or *TAG*, that provides a rich set of tools for viewing and editing annotations that capture complex natural language phenomena, such as deeply nested event structures (Figs. 1 and 6), syntactic dependencies (bottom portion of Fig. 1), undirected relations, coreference resolution, and morphological parses (Fig. 5). *TAG* provides easily readable output for complex events and can be used to illustrate non-projective representations and even link predicate-less relations. It supports multiple languages and data formats, allows users to view layers of annotations simultaneously, and includes interactive features for filtering and editing annotations. Moreover, *TAG* implements progressive rendering or text annotations, making it possible to load in large documents without loss in rendering speed.

TAG runs on any modern web browser and is available as open source software from <https://github.com/CreativeCodingLab/TextAnnotationGraphs>.

2. Related Work

TAG is inspired by the *BRAT* rapid annotation tool (Stenetorp et al., 2012). *BRAT* is widely used for representing syntactic structure, but can also represent semantic events, and has been applied to a range of domain-specific NLP tasks, including biomedical data (Verspoor et al., 2013). *BRAT* supports a range of useful features that improve the overall efficiency of manual annotation tasks. However, *BRAT* does not support the ability to draw links between links, which makes it difficult to represent relations linking several predicate-less relations, a feature necessary to completely describe complex events. Complex relations are previously explored by the authors in a range of different visualization projects that represent hierarchically-nested and/or clustered data derived from the machine reading of scientific texts describing biochemical events (Dang et al., 2015; Dang et al., 2016; Dang and Forbes, 2017; Dang et al., 2017; Forbes et al., 2018; Murray et al., 2017; Paduano and Forbes, 2015; Paduano et al., 2016).

In addition to *BRAT*, a range of newer projects investigate visual encodings for specific annotation tasks. For example, Sarnat et al. (2017) introduce a web interface for exploring a parse tree. By entering in a sentence, an interactive visualization is created that includes expand/collapse functionality, positional and color cues, explicit visual support for sequential structure, and dynamic highlighting to convey node-to-text correspondence. The tool includes an unusual representation of sequential structure as a container of linked nodes. While this may help a user to distinguish relevant structural elements, it demands a large portion of the screen, which can make it difficult to view relationships over longer sequences of text. *TAG* also utilizes a range of visual encodings to identify relationships and types, and includes an alternative representation of linguistic relationships. However, in *TAG*, a user can fluidly switch between these representations or view them side by side.

Another representation of a parse tree is used in the *displaCy* software, introduced by Montani (2017). *displaCy*, however, is not equipped for annotation tasks and does not provide any interaction beyond the ability to scroll through a sentence (although more sophisticated functionality could potentially be built on top of it). By default, the entire body of text is placed on a single row, making it unwieldy for representing longer swaths of text.

WebAnno is a flexible tool that supports multiple annotation layers, and includes features to facilitate quality control, annotator management, and curation (Yimam et al., 2014; de Castilho et al., 2016). The visual representation is similar to *BRAT*, and the interface focuses mainly on resolving disagreeing annotations between users. *WebAnno* includes a variety of built-in annotation layers, such as dependency relations, co-reference chains, and lemma forms, but in *WebAnno* annotation expressiveness is limited in that it is not possible to create nested arcs that link to other arcs.

Purgina and Mozgovoy (2017) introduce *WordBricks*,

which also utilizes a container layout to identify linguistic structure while maintaining an explicit sequential representation that is easy to read. Although their approach eschews links that may lead to visual clutter, the *WordBricks* layout can be difficult to read once the container is more than even a few levels deep, or in cases where the user wishes to quickly identify relationships between words or phrases across sentences or within longer passages.

3. Text Annotation Graphs (TAG)

In this section, we introduce the main visualization and interaction features of the *TAG* software application. It consists of three panels: the annotation panel, the tree panel, and the options panel. Upon opening the software via any modern browser, *TAG* shows the user a main annotation panel consisting of an example text snippet with both semantic and syntactic annotations. A secondary tree panel shows a “semantic summary” of a selected subgraph, in which text is organized using a network layout that emphasizes event relationships rather than text sequences. This alternative representation provides an overview of the event that can augment the more traditional reading of a parse tree. Additionally, it can be used to highlight the results of an event extraction process, and used as input for a structured search across other annotated documents. A third panel can be displayed on demand that provides the user with a range of options to change the visual appearance of text, to filter out particular annotation types, and to toggle on or off syntactic or semantic annotations.

3.1. Loading Data

TAG is meant to be agnostic to data schema and currently supports the following formats: *BRAT* standoff (Stenetorp et al., 2012), *CoNLL-X* (Buchholz and Marsi, 2006), and *bio-C* (Comeau et al., 2013). Additional formats can be added as desired by following the import templates for the supported formats. The software looks for files in a special ‘data’ folder, and populates a drop down list to provide the option for the user to load in any of those data files. By choosing an item in this drop down list, which is located in the top left of the page, the data file is loaded into the *TAG* application, and populates the annotation panel. A taxonomy file can also be associated with a data file, which enables *TAG* to color annotations based on their taxonomic type, and makes it possible for users to search and filter by type.

3.2. Annotation Panel

The annotation panel takes up the majority of the *TAG* application, filling up the entire page, or only the top two thirds of the page if the tree panel is opened. The annotation panel presents the text from the data file across many rows, requiring the user to scroll down to see text from longer passages. *TAG* uses a progressive rendering strategy so that it is possible to display even large documents without loss in rendering speed. In each row, an annotation graph is shown both above and below the text. At the bottom, syntactic annotation is shown; At the top, semantic annotation is shown. Either syntactic or semantic annotations can be toggled on or off on demand.

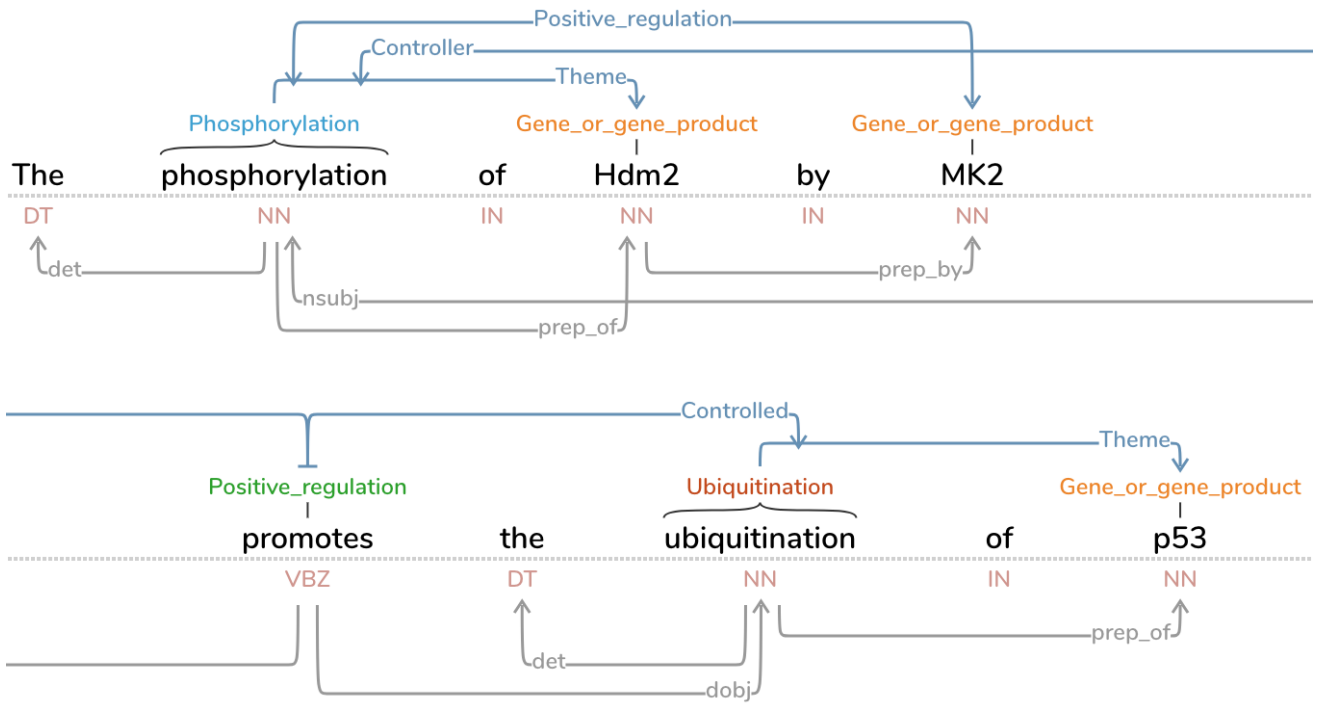


Figure 2: This screen capture shows the layout of a single sentence where annotation data is linked between elements which are presented on different rows. A user can interactively reposition words and links as desired across multiple rows in order to emphasize particular sections of text.

Figure 3: This screen capture shows a user updating the color of all nodes of type 'Gene_or_gene_product' via the interactive color picker. In the background, we can see that all examples of this annotation are updated dynamically. Selecting an annotation type that is a parent of 'Gene_or_gene_product', such as 'MacroMolecule' or 'Entity', would also update the color of 'Gene_or_gene_product', along with all its other children.

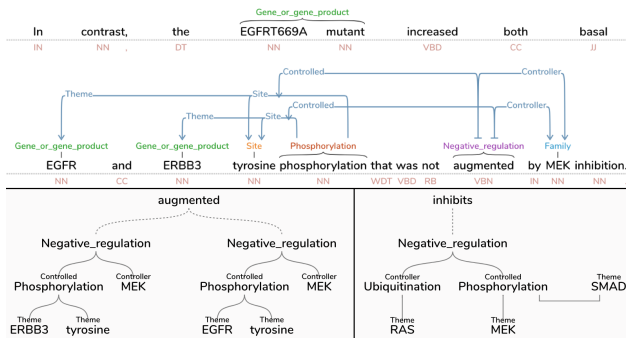


Figure 4: The tree panel shows a semantic summary of a selected annotation. The top and bottom left show the same annotation in both formats, in the tree panel emphasizes the negative regulation events triggered by the word ‘augmented.’ The bottom right shows the semantic summary of the sentence shown in Fig. 1, highlighting the events triggered by the word ‘inhibits.’

TAG supports trigger-free relation mentions, providing the ability for the annotation graphs to utilize uni- or bi-directional links between nodes, links, or a combination of nodes and links. The layout of the links is organized using a custom algorithm that makes text labels on the graph easy to read, and minimizes edge crossings. Fig. 1 shows an example of data loaded into the annotation panel that requires only one row.

A user can interactively reposition the words in the text, including seamlessly moving words to different rows. This allows the user to emphasize particular word sequences by clustering a series of words closer together. Similarly, the vertices of links can be repositioned to improve the readability of the annotation graph. Fig. 2 show a view of an annotated sentence across two rows. Initially, it was positioned by default along a single row; here we see an example of how a user might choose to arrange the sentence interactively. In this case, annotation data is linked between elements which are on different rows and the syntax tree has been toggled on.

Annotations can be edited on demand. By clicking on a node, the user can select a type from a provided taxonomy, or freely type to replace the node’s current annotation. Links can be repositioned so that one or more of the connections can instead connect to different link or node. A node or link can also be hidden or deleted, which also hides or removes any links directly connected to it.

Additionally, a user can change the color of a single node (or all nodes of a particular type) by clicking on a node, which pops up an option panel, which allows the user to select a color dynamically. Figure 3 shows a user updating all nodes of type ‘Gene_or_gene_product’ to be colored red. When changing the color of a type, a user can also update all of its subtypes as well, if desired.

3.3. Tree Panel

By default the annotation panel fills up the entire screen, but the bottom third of the application can be used to provide an alternative representation of the annotated text. By double-clicking on any word or annotation, the tree panel

appears in the application, displaying a tree whose root node is the element that was clicked on. The tree representation makes it easier to see the structure of the relations between annotation events, which may be more difficult to read in the annotation panel, especially when related events are spread across multiple sentences or rows, or when a single sequence of text contains multiple events. Of course, the tree representation removes the context from which the events are extracted. However, by presenting both views at once, the user is able to see the summary semantics of the annotation relations alongside the annotations embedded in the text.

Fig. 4 shows an example of the tree panel. On the top, the annotation panel shows a sentence containing two negative regulation events whose representation has a somewhat dense layout, even after the user has repositioned the event so that all involved words are on the same row. On the bottom left, we see the simplified tree representation of the events, which appears when the user clicks on the word “augments.” On the bottom right, we see another tree representation—the sentence described in Fig. 1—that highlights the annotations related to the word “inhibits,” and which shows an example of a sibling relationship.

3.4. Option Panel

The option panel provides various ways for the user to customize the software. The options include: the ability to toggle on or off the annotations above or below the text; optimization options, such as the ability to hide the annotation graph when repositioning words; and a mechanism for filtering annotations that contain (or do not contain) a specified taxonomic type. The user can also change the color of types defined in the taxonomy to highlight particular elements (a biomedical taxonomy can be seen in Fig. 3).

4. Examples

Although *TAG* can be used for a wide range of annotation tasks, here we highlight two examples that showcase situations that are difficult to represent using other software.

4.1. Morphological Parse

Fig. 5 shows a pair of examples that demonstrate semantic differences arising from morphological derivations of “unlockable,” which, as noted by Larson and Ludlow (1993), could mean either “cannot be locked” (top) or “can be unlocked” (bottom). As we can see, trees such as these cannot easily be visualized in *BRAT*, but can be represented in *TAG*.

4.2. Event Annotation

Fig. 6 shows a comparison of how *BRAT* and *TAG* handle the output of Reach (Valenzuela-Escárcega et al., 2015) when processing the title text of the paper “Induction of p21 by p53 following DNA damage inhibits both Cdk4 and Cdk2” (He et al., 2005). Looking at extracted event E1, we can see that the Reach system finds that the positive activation event on the word “Induction” indicates p53 as the controller for p21. In the *BRAT* representation, the “controller” and “controlled” relations are separated, whereas in *TAG* it is easier to see that “induction” links p21 and p53




Figure 5: Two interpretations for “unlockable” arising from differences in morphological derivation displayed in both *BRAT* (left) and *TAG* (right). The top parse models the interpretation “cannot be locked”, while the bottom shows “can be unlocked”.

through a single positive activation trigger. Similarly, *TAG* is able to recognize that the a single word serves as a trigger for two negative regulation events, and presents it as a single event with two simultaneous relations. Here, it is the entire positive activation event (E1) that serves as the controller for the inhibition of both Cdk4 and Cdk2.

5. Conclusion

This paper described the main functionality of the *Text Annotation Graph* software, showcasing its increased expressivity for visually representing complex natural language phenomena for even large text documents. The software has been used so far to successfully represent events extracted from biomedical articles, but we expect that it could prove useful for annotation tasks across a range of domains. Ongoing work is focused on extending *TAG* to support comparative tasks, for instance, to use both the top and bottom half of each row to present the semantic annotations extracted from different systems. We are also investigating the use of *TAG* to provide a human readable visualization of an active learning loop, enabled through saving “diffs” between the loaded data and the edited data. Since *TAG* makes it easy to read and edit annotated events, we are exploring its use as a front end for crowdsourced user studies, such as those administered via Mechanical Turk, in order to gather data that can be used to train classifiers.

TAG is written in JavaScript and runs in any browser. It makes extensive use of the SVG.js library,¹ which produces PDF quality output that is suitable for embedding in both online and printed documents. The software is open source and freely available from our GitHub code repository.²

Acknowledgements

This work was funded by the DARPA Big Mechanism program under ARO contract W911NF-14-1-0395 and by

¹<http://svgjs.com/>

²<https://github.com/CreativeCodingLab/TextAnnotationGraphs>

T1	Gene_or_gene_product	13	16	p21
T2	Gene_or_gene_product	20	23	p53
T3	BioProcess	34	44	DNA damage
T4	Gene_or_gene_product	59	63	Cdk4
T5	Gene_or_gene_product	68	72	Cdk2
T6	Positive_activation	0	9	Induction
T7	Negative_regulation	45	53	inhibits
E1	Positive_activation::T6	Controller::T2	Controlled::T1	
E2	Negative_regulation::T7	Controller::E1	Controlled::T4	
E3	Negative_regulation::T7	Controller::E1	Controlled::T5	

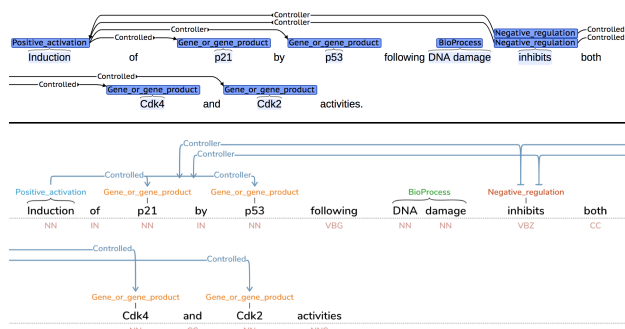


Figure 6: Here we see screen captures of the same Reach events (top) in both *BRAT* (middle) and *TAG* (bottom). In *TAG*, it is clear that E2 and E3 both use E1 as a controller of Cdk4 and Cdk2.

the Bill and Melinda Gates Foundation HBGDKi Initiative. Gus Hahn-Powell, Marco Valenzuela-Escárcega, and Mihai Surdeanu declare a financial interest in lum.ai. This interest has been properly disclosed to the University of Arizona Institutional Review Committee and is managed in accordance with its conflict of interest policies.

Bibliographical References

- Björne, J., Ginter, F., Pyysalo, S., Tsujii, J., and Salakoski, T. (2010). Complex event extraction at PubMed scale. *Bioinformatics*, 26(12):i382–i390.
- Buchholz, S. and Marsi, E. (2006). CoNLL-X shared task on multilingual dependency parsing. In *Proceedings of the Tenth Conference on Computational Natural Language Learning*, pages 149–164.
- Comeau, D. C., Islamaj Doğan, R., Ciccarese, P., Cohen, K. B., Krallinger, M., Leitner, F., Lu, Z., Peng, Y., Rinaldi, F., Torii, M., et al. (2013). BioC: A minimalist approach to interoperability for biomedical text processing. *Database*, 2013:bat064.
- Dang, T. and Forbes, A. G. (2017). CactusTree: A tree drawing approach for hierarchical edge bundling. In *Proceedings of the 10th IEEE Pacific Visualization Symposium (PacificVis)*, pages 210–214, Seoul, Korea, April.
- Dang, T., Murray, P., Aurisano, J., and Forbes, A. G. (2015). ReactionFlow: An interactive visualization tool for causality analysis in biological pathways. *BMC Proceedings*, 9(6):S6, August.
- Dang, T., Pendar, N., and Forbes, A. G. (2016). TimeArcs: Visualizing fluctuations in dynamic networks. *Computer Graphics Forum*, 35(3):61–69, June.
- Dang, T., Murray, P., and Forbes, A. G. (2017). BioLinker: Bottom-up exploration of protein interaction networks. In *Proceedings of the 10th IEEE Pacific Visualization*

- Symposium (PacificVis)*, pages 265–269, Seoul, Korea, April.
- de Castilho, R. E., Mújdricza-Maydt, E., Yimam, S., Hartmann, S., Gurevych, I., Frank, A., and Biemann, C. (2016). A web-based tool for the integrated annotation of semantic and syntactic structures. In *Proceedings of the LT4DH workshop at COLING 2016*.
- Forbes, A. G., Burks, A., Lee, K., Li, X., Boutillier, P., Krivine, J., and Fontana, W. (2018). Dynamic influence networks for rule-based models. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):184–194, January.
- He, G., Siddik, Z. H., Huang, Z., Wang, R., Koomen, J., Kobayashi, R., Khokhar, A. R., and Kuang, J. (2005). Induction of p21 by p53 following DNA damage inhibits both Cdk4 and Cdk2 activities. *Oncogene*, 24(18):2929.
- Lander, A. D. (2010). The edges of understanding. *BMC Biology*, 8(1):40.
- Larson, R. K. and Ludlow, P. (1993). Interpreted logical forms. *Synthese*, 95(3):305–355.
- Montani, I. (2017). displaCy Dependency Visualizer. <https://github.com/explosion/displacy>, accessed Sept. 24, 2017.
- Murray, P., McGee, F., and Forbes, A. G. (2017). A taxonomy of visualization tasks for the analysis of biological pathway data. *BMC Bioinformatics*, 18(2):21–1–13.
- Paduano, F. and Forbes, A. G. (2015). Extended LineSets: A visualization technique for the interactive inspection of biological pathways. *BMC Proceedings*, 9(6):S4, August.
- Paduano, F., Etemadpour, R., and Forbes, A. G. (2016). BranchingSets: Interactively visualizing categories on node-link diagrams. In *Proceedings of the 9th International Symposium on Visual Information Communication and Interaction (VINCI)*, pages 9–16, Dallas, Texas, September.
- Purgina, M. and Mozgovoy, M. (2017). Visualizing sentence parse trees with wordbricks. In *Proceedings of the 3rd IEEE International Conference on Cybernetics (CYBCON)*, pages 1–4.
- Sarnat, A., Joshi, V., Petrescu-Prahova, C., Herrasti, A., Stilson, B., and Hopkins, M. (2017). Interactive visualization for linguistic structure. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Stenetorp, P., Pyysalo, S., Topić, G., Ohta, T., Ananiadou, S., and Tsujii, J. (2012). BRAT: A web-based tool for NLP-assisted text annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107.
- Valenzuela-Escárcega, M. A., Hahn-Powell, G., Hicks, T., and Surdeanu, M. (2015). A domain-independent rule-based framework for event extraction. In *Proceedings of ACL-IJCNLP 2015 System Demonstrations*, pages 127–132.
- Verspoor, K., Jimeno Yepes, A., Cavedon, L., McIntosh, T., Herten-Crabb, A., Thomas, Z., and Plazzer, J.-P. (2013). Annotating the biomedical literature for the human variome. *Database*, 2013:bat019.
- Yimam, S. M., Biemann, C., de Castilho, R. E., and Gurevych, I. (2014). Automatic annotation suggestions and custom annotation layers in webanno. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 91–96.